UCRL-JC-120319
PREPRINT

Multispectral Image Fusion for Detecting Land Mines

Gregory A. Clark, Sailes K. Sengupta, William D. Aimonetti, Frank Roeske, John G. Donetti, David J. Fields, Robert J. Sherwood, and Paul Schaich

This paper was prepared for submittal to the SPIE Orlando Symposium Orlando, Florida April 17–21, 1995

April 1995

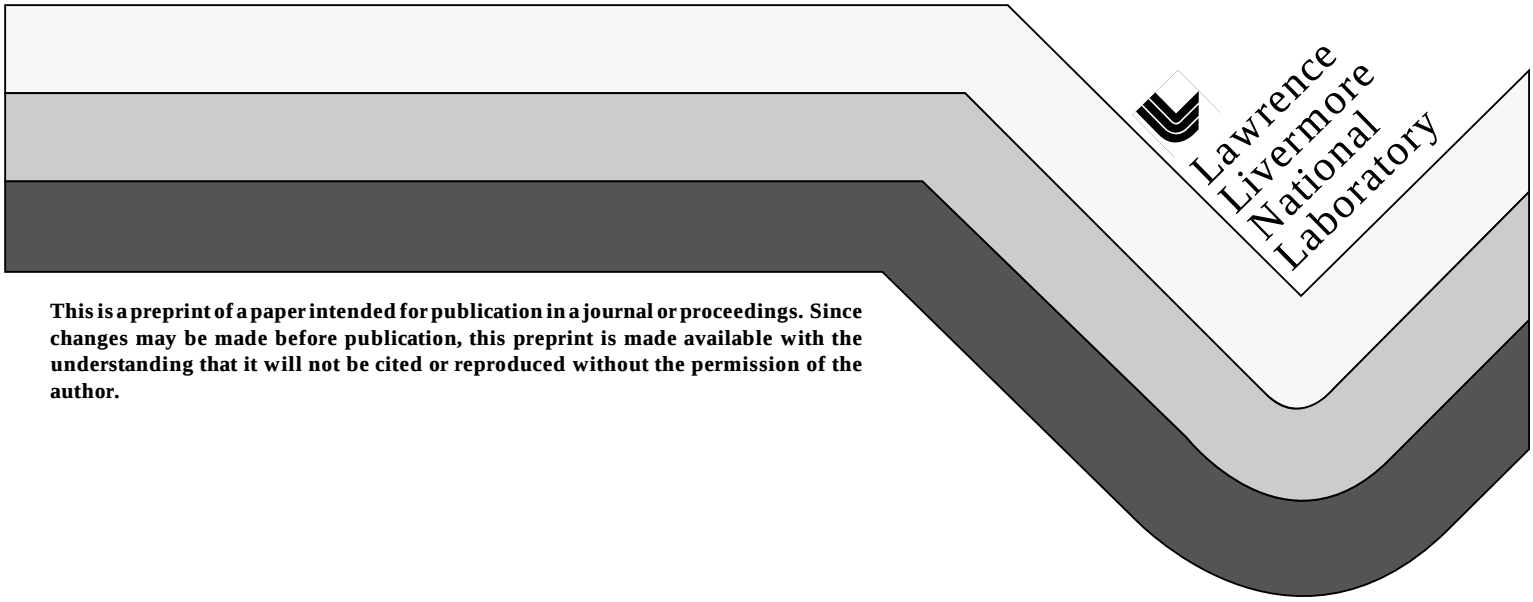

Multispectral Image Fusion for Detecting Land Mines

Gregory A. Clark, Sailes K. Sengupta, William D. Aimonetti,
Frank Roeske, John G. Donetti, David J. Fields, Robert J. Sherwood, Paul C. Schaich

Lawrence Livermore National Laboratory
7000 East Avenue, L-156, Livermore, CA 94550

ABSTRACT

Our system fuses information contained in registered images from multiple sensors to reduce the effects of clutter and improve the ability to detect surface and buried land mines. The sensor suite currently consists of a camera that acquires images in six bands (400nm, 500nm, 600nm, 700nm, 800nm and 900nm). Past research has shown that it is extremely difficult to distinguish land mines from background clutter in images obtained from a single sensor. It is hypothesized, however, that information fused from a suite of various sensors is likely to provide better detection reliability, because the suite of sensors detects a variety of physical properties that are more separable in feature space. The materials surrounding the mines can include natural materials (soil, rocks, foliage, water, etc.) and some artifacts.

We use a supervised learning pattern recognition approach to detecting the metal and plastic land mines. The overall process consists of four main parts: Preprocessing, feature extraction, feature selection, and classification. These parts are used in a two step process to classify a subimage. The *first step*, referred to as feature analysis, determines the features of sub-images which result in the greatest separability between the two classes, "mine" and "background." The *second step*, image labeling, uses the selected features and the decisions from a pattern classifier to label the regions in the image which are likely to correspond to mines.

We extract features from the images, and use feature selection algorithms to select only the most important features according to their contribution to correct detections. This allows us to save computational complexity and determine which of the spectral bands add value to the detection system. The most important features from the various sensors are fused using a supervised learning pattern classifier (the probabilistic neural network). We present results of experiments to detect land mines from real data collected from an airborne platform, and evaluate the usefulness of fusing feature information from multiple spectral bands. We show that even with preliminary data and limited testing, the performance (specified in terms of probability of detection and probability of false alarm) is very promising. The novelty of the work lies mostly in the combination of the algorithms and their application to the very important and currently unsolved operational problem of detecting minefields from an airborne standoff platform.

1. INTRODUCTION

The goal of this work is to detect and locate buried and surface land mines, given multiple registered images of regions of the earth obtained from a suite of visible wavelength sensors. Past research has shown that it is extremely difficult to distinguish objects of interest from background clutter in images obtained from a single sensor. It has been hypothesized, however, that information fused from a suite of various sensors is likely to provide better detection reliability, because the suite of sensors measures a

variety of physical properties that are more separable in feature space. Materials surrounding the objects of interest can include natural materials (soil, rocks, foliage, water, etc.) and artifacts (objects made of metal, plastic and other materials).

The sensor suite for this work consists of a filter wheel camera that acquires images in six bands (400nm, 500nm, 600nm, 700nm, 800nm and 900nm). The detection system uses advanced algorithms from the areas of automatic target recognition (ATR), computer vision, signal and image processing, and information fusion. The system uses both physical principles and image processing for image interpretation.

This work is application research in progress. The individual algorithms used are advanced, but mostly known, and the novelty of the work lies in the combination of the algorithms and their application to the very difficult and important problem of detecting buried land mines. To date, no successful operational system exists for airborne standoff detection of buried land mines. At the current time, our data set is limited, in that we have a small sample size.

2. EXPERIMENTS AND MEASUREMENTS

The images were acquired using an airborne platform and a six-band visible wavelength camera. The raw data images have size 720 pixels x 480 pixels in six bands (400nm, 500nm, 600nm, 700nm, 800nm and 900nm). The images from the six bands are coregistered. The targets (mines) are mostly of the surface type, but the experiment included some buried mines. There is a mixture of metal mines, plastic mines, and mine surrogates (concrete stepping stones). The sizes and shapes of the mine targets vary. Shapes include circular, near circular, rectangular, tubular (long and narrow), and square. The size of the largest dimension of the targets varies from about three inches to about two feet. Artifacts in the images include intentionally-placed resolution panels. Other natural objects in the images include foliage, soil, rocks, water, etc.

3. DATA FUSION AND AUTOMATIC TARGET RECOGNITION (ATR)

This section provides an overview of the general ATR algorithms and philosophies behind them. Section 4 describes the details of how the algorithms are applied.

3.1 Supervised learning

The ATR system is a supervised learning classifier for which we define two classes; “mine” and “background” (not mine). Therefore, the ATR system is designed to classify image regions as either mine regions or background regions. The supervised learning approach is applied in two steps; training and testing.

3.1.1 Training

In the *training step*, we present the classifier with a “training set” of examples (sub-images, or “tiles”) of mine and background regions, along with their associated “ground truth,” or prior knowledge of the true class to which each example belongs (mine or background). Once the classifier is trained to successfully classify the training data with acceptable performance measured by probability of detection and probability of false alarm [4, 8-16], we move to the testing step.

3.1.2 Testing

The *testing step* consists of using the trained classifier to process an image that was not included in the training set and making the appropriate classifications. For this application, the testing occurs in two very different ways: testing with "tiles" (image samples or subimages), and testing an entire image.

3.1.2.1 Testing with tiles

This means that we save aside image tiles that were not used for training and apply them to the trained classifier. A confusion matrix, including calculations of probability of detection and probability of false alarm is generated to evaluate performance.

3.1.2.2 Testing with an entire image (image labeling

An analysis window of the same size as the training tiles is raster-scanned over the image. A new labeled image is constructed as follows. At each pixel in the testing image, features are calculated for the pixels in the analysis window, and the classifier classifies the center pixel in the window as belonging to either the class "mine" (this pixel in the labeled image is assigned a value of 1) or the class "background" (this pixel in the labeled image is assigned a value of 0). The resulting binary labeled image contains only ones representing mine pixels and zeros representing background pixels. Mine regions appear in the labeled image. A post-processing step (described later in this paper) is then used to apply size and shape constraints to the detected mine regions.

3.1.3 The Hold-One-Out Method of Supervised Learning [4, 8-15, 29]

The ideal supervised learning paradigm involves having a large set of N data samples available which are divided using an empirical rule of thumb into a training subset (about 2N/3 samples) and a testing subset (about N/3 samples). However, when the number of available samples, N, is small, we can only approximate this ideal case. A well-known and accepted approximation is called the "hold-one-out" method. Here, we start by using all of the N available data samples, except for one which is "held out," to train the classifier, and test the one held out sample. Next, we insert the held out sample back into the training set and hold out another sample for testing. We repeat the procedure, holding out one sample and training with the remaining samples at each iteration until all N of the samples have been held out once.

For our problem, we can interpret and use the hold-one-out method in either or both of two ways; (1) hold one mine or background sample out, and/or (2) hold one image out. We use both techniques. We designate most of the images for training and designate the remaining images for labeling. We use the hold-one mine or background sample out method for training using samples derived from the training images. Then, we test using the images held out for testing only.

3.1.4 ATR processing

The overall target recognition process is depicted in Fig. 1 and consists of four main parts: preprocessing, feature extraction, feature selection, and classification. These parts are used in a two step process to classify a subimage. *The first step*, referred to as feature analysis, determines the features of sub-images which result in the greatest separability among the classes. The *second step,* image

labeling, uses the selected features and the decisions from a pattern classifier to label the regions in the image which are likely to correspond to buried mines.

3.2 Image Preprocessing

3.2.1 Image Cropping

The effective area of the images available for automatic detection is limited due to camera time stamps on the images. These time stamps obscure some of the mines and make it necessary for us to crop the images. The cropped image uses columns 21 to 695 and rows 110 to 398 to assure that we do not have any of the text in our field of view for detection. This limits the number of targets we can use in our

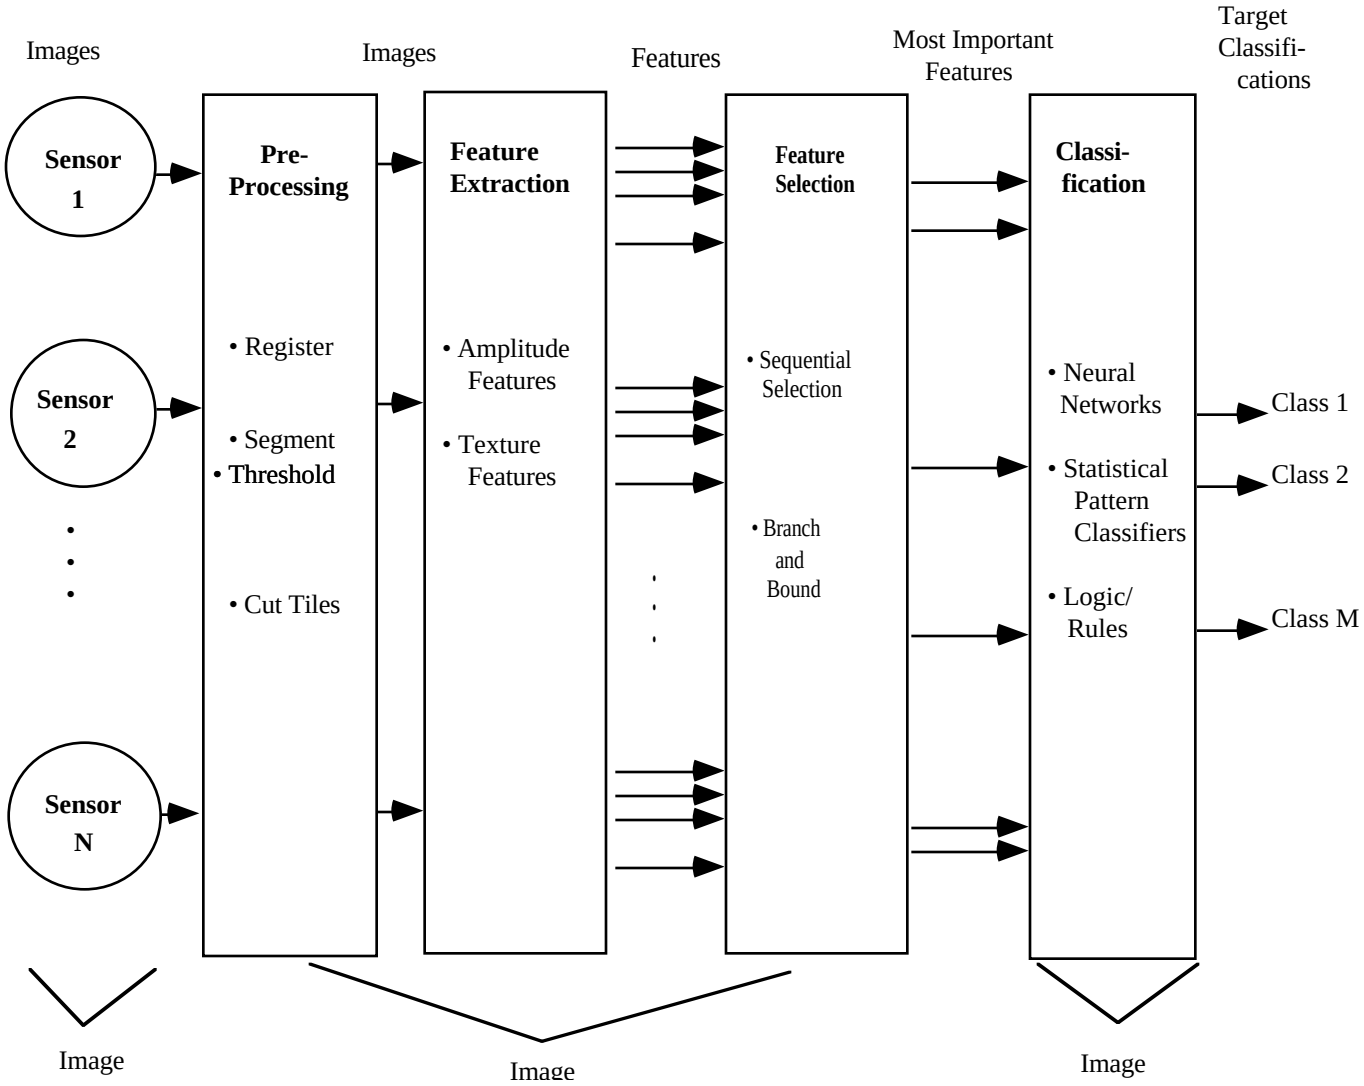

Figure 1. Fusion/Automatic Target Recognition Depend Heavily Upon Proper Image Representation.

training set for the classifier to 133 out of a total of 209 possible objects listed in the ground truth table. The time stamps eliminated 36% of the possible targets.

3.2.2 Normalization:

The images are normalized with respect to the background by subtracting the mean of the background from the images and dividing this result by the standard deviation of the background. This normalization aids in computing some of the image features and it makes the classifier less sensitive to absolute units, which can vary greatly with physical properties of the site from image to image.

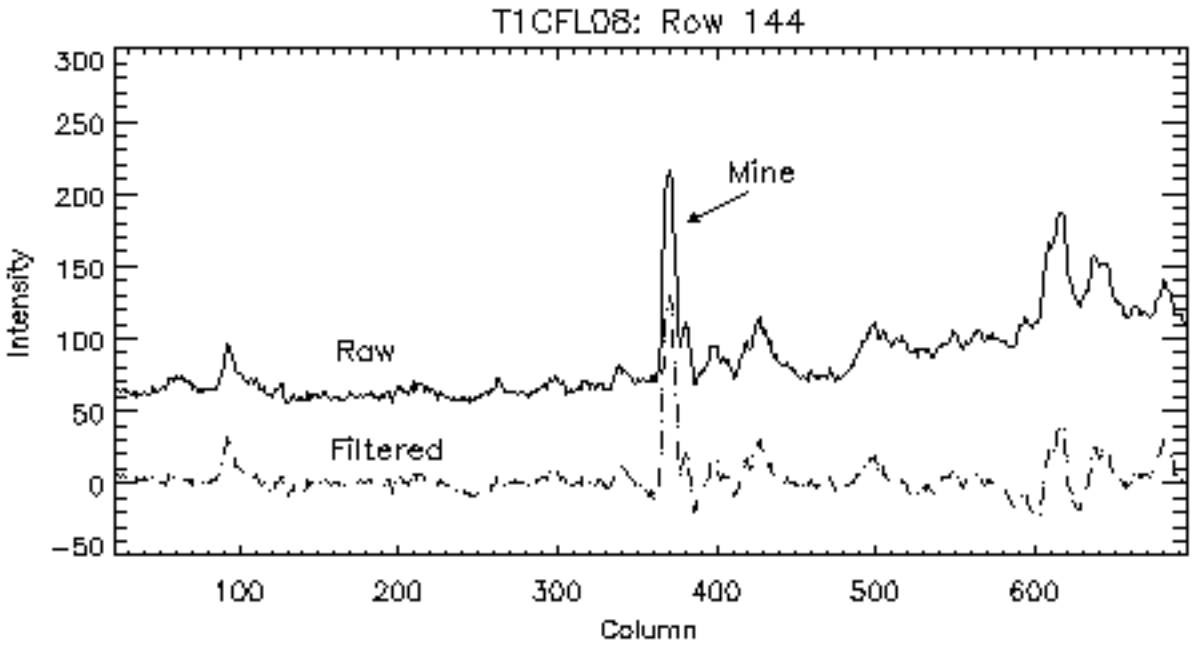

Figure 2. These row lineouts show and example of the trend in the raw data and the removal of the trend in the filtered data.

3.2.3 Trend Removal:

The data also contain a spatial trend in the background. Background values across the image differ by up to a factor of two from one side to the other (fig1). As a result of this the signal-to-noise ratio (SNR) varies across the image and, in some cases, it was very poor. Varying SNR's cause the training sets to become less reliable and, of course, where the SNR is low, we observe a decrease in contrast between mines and backgrounds. This is a problem with the camera, and steps are being taken to correct it. Until the corrections are made, however, we must cope with the data as it is.

We high-pass filtered the cropped data sets to reduce this trend. For each pixel, the raw value was replaced by the difference of the raw value and the average of a 30 by 30 area surrounding the pixel. This operation greatly reduced the background trend in the raw image (see fig 2).

3.2.4 Tile Cutting

We use ground truth information about the scene to cut out N x N pixel tiles (sub-images) centered around mine regions and background regions. These tiles become the training samples used for pattern classification. The specifics are provided in section 4.

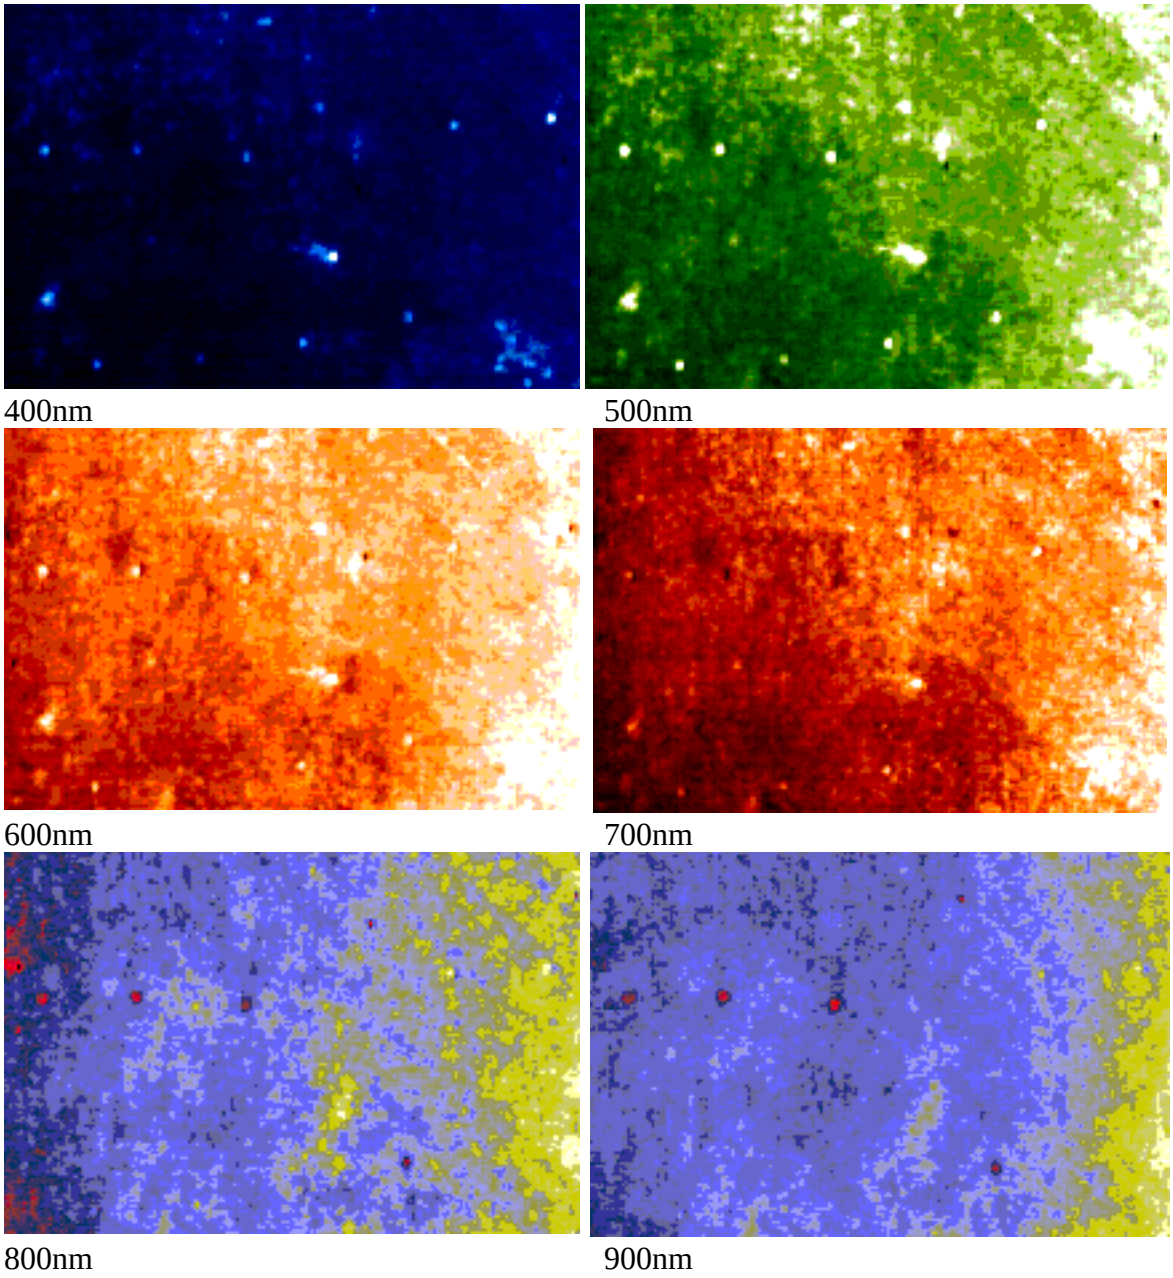

Figure 3: These figures show the data set T1CFL08 after highpass filtering for trend removal.

3.3 Feature Extraction

As part of the supervised learning paradigm, we create a set of training samples and a set of test samples. Each image is divided into N x N pixel subregions centered at mine pixels and background pixels. For this data set, N = 5. We create a "tile stack" of six subimages or "tiles," one for each frequency band, so the 3D stack is of size 5 pixels X 5 pixels X 6 pixels. We chose the tile size so as to be contained entirely inside the boundary of the mine image, and not contain a significant amount of background information.

During the training phase, the spatial locations of the mine tiles are given by the ground truth provided by the Coastal System Station (CSS). The background training set is obtained by cutting tiles of the same size as the mine tiles. The locations chosen for cutting background tiles, however, are randomly chosen from among the possible background areas in the scene.

Given preprocessed sub-images, we compute a vector of statistical features from the pixel values in the sub-images. Typical features include amplitude histogram features and texture features [2, 3]. Our philosophy is to use the simplest features that are effective, so we currently use only the amplitude histogram features (mean, standard deviation, skewness, kurtosis, energy and entropy). Amplitude features are statistical moments of the probability density function (pdf) for the pixel intensity values within the tiles. Sample moments are used to estimate the features. Future plans involve testing the effectiveness of features such as texture features that exploit spatial information.

For these data, the size of the training set for mines varies from about 100 tile stacks (samples) for the case in which we included all types of targets, to about 50 samples (tile stacks) for the case in which we included only metal mines as targets. The training set for the background class contains about twice as many samples as the mine class in all cases.

3.4 Feature Selection

Human experts generally classify objects based on a very few of the most important attributes in the image. The fundamental function of the feature selection process is to select the most useful information from the representation vector and present it in the form of a relatively low-dimensional pattern vector removing any redundant and irrelevant information which may have a detrimental effect on the performance of the classifier. A useful by-product in the process is knowledge about the discriminatory potential of the features and the associated highest achievable performance for a given set of features. Statistical decision theory tells us that the probability of misclassification is a decreasing function of the number of features provided, if the sample size is very large. In practice however, only a small number of training sets is available and estimation errors are no longer negligible. Since the number of parameters and the associated estimation errors increase rapidly with dimension, it may be advantageous to sacrifice some useful information in order to keep the number of these parameters to a minimum.

An important goal in our work is to use feature selection techniques to choose the subset of features that contribute most to correct classification. We gain two main benefits from this approach. First, we wish to *minimize the computational complexity* of our processing algorithms, so they can eventually be implemented in "real time." Second, we wish to *determine which sensors are the most important* for classification. By rank ordering the features according to their importance for classification, we are able to eliminate from consideration sensors which do not contribute significantly. Feature selection is typically accomplished by computing a distance measure which is the sum of probabilistic distances between all pair-wise combinations of classes [3,4]. Commonly used algorithms include branch and bound, sequential forward selection, and sequential backward selection [3,4]. For this study, we used the sequential forward selection algorithm because it is very effective and computationally efficient.

We must pay careful attention to an important relationship between the number of features used and the sample size (in this case, the number of independent tiles) in the training set. A combination of theoretical and empirical studies has led to the following *rule of thumb* [4, 21]:

No. of independent training samples needed per class ≥ 5 (No. of Features)

For example, if we have a feature vector of dimension 10, then we need at least 50 training samples in each class to support the classifier. In fact, many researchers recommend using many more than five times the dimension of the feature vector.

This rule of thumb has been validated and has proved to be of great value in mine detection and a variety of other applications the authors have studied [8-16]. The theoretical reasoning for the rule of thumb is based upon the fact that covariance matrices are used in feature space class separability measures and in many classification algorithms. The rule of thumb reflects the number of training samples required to ensure in practice that the covariance matrix is estimated with sufficient precision.

An important implication of this rule of thumb is an *upper bound on the number of features to use*, given the number of independent training samples. Note that if the sample size is small, as it is in this mine detection study, it severely limits the number of features we can use. In our work, for example, we were limited to about 3 features, because our small sample size would not support more samples. This is discussed in greater detail in the section on processing and results.

3.5 Classification

We choose to use the probabilistic neural network (PNN) as the classifier, for reasons described in [6]. The PNN is a Bayesian classifier based upon the Parzen estimator of conditional probability density functions (pdf) [6]. The PNN has the desirable property that it provides the Bayes optimal pdf estimate in the limit as the number of training samples approaches infinity. For our problem, given a feature-vector X as input data, the PNN calculates the values f(X|mine) and f(X|background). These pdf values can be used to calculate the posterior probability of the source given X, P(mine|X), and the posterior probability of the background given X, P(background|X). Classification of the vector X is obtained by selecting the class with the largest value of the posterior probabilities given above.

Because we have a small sample size, we use the "hold one out" method for training and testing, as described in section 3.1.

3.6 Image Labeling

Once the classifier is trained, it is used to analyze an image not included in the training sets. An analysis window of the same size as the training tiles is raster-scanned over the image. At each pixel in the image, features are calculated for the pixels in the analysis window, and the PNN classifies the center pixel in the window as belonging to either the class "mine" or the class "background." The resulting binary image (containing only ones representing mine pixels and zeros representing background pixels) is called the "labeled image." This labeled image provides us with an indication of the locations of probable mine pixels. The labeled image often contains "false alarm" pixels where the PNN classified the pixel as a mine pixel, when it was in fact a background pixel. When a large number of false alarms occur, their number can be greatly reduced by postprocessing the labeled image as described next.

3.7 Image Post-processing

After the labeling step, we use region-based operations to the labeled image to automatically identify regions, and apply size and shape constraints which can help eliminate false alarm regions. We first apply a morphological operator to the labeled image. We use a 3x3 pixel structuring element to successively erode the labeled image, then dilate the eroded image. This operation serves to eliminate

many of the small false alarms from the labeled image. The operation of erosion followed by dilation is called an 'opening. The opening sieves out objects that are smaller than the structuring element, but avoids a general shrinking of the image [30, 34].

Since we know the physical size of our targets, we can apply size constraints to eliminate from the opened image objects that are too large or too small to represent targets. First we perform a connected components analysis on the opened image [30, 34]. This operation assigns all adjacent pixels that form a region a unique number (index). We then eliminate detected object regions that are too large or too small according to prior knowledge we have about the mine size from ground truth. Given the true mine size, we discard detected regions that are smaller than 66% of the size the region is supposed to be, according to ground truth for a mine. The number "66%" was chosen based upon knowledge that the labeling process tends to erode the size of the region detected, and by some experimentation with the data set. In future work, this number can be better tuned to the data with more effort, or more sophisticated algorithms can be employed. The important concept to note is that the use of size constraints is a very useful tool for eliminating small false alarms and for using prior knowledge in the analysis.

4. PROCESSING RESULTS AND DISCUSSION

4.1 Training Data Set

Several training sets were generated to accommodate the intended targets to be detected. These sets contained various types of targets : All targets, all targets without resolution panels, plastic and metal mines only, metal mines only, and plastic mines only. We searched a space of 54 features (nine features for each band) for three features selected by the Sequential Forward Selection algorithm [4]. Selected features differed for each type of training table. The Probabilistic Neural Net (PNN) was tuned for optimal probability of correct classification by adjusting the smoothing parameter, sigma [6]. The smoothing parameter defines the width of the Parzen window used in the PNN. The table below summarizes the type of training table, the selected features, and the optimal sigma that was obtained in each case.

training table:	band/features:	sigma:
1. All targets	400nm Max, 400nm Energy, 400nm Entropy	0.09
2. All targets, no Res panels	400nm Max, 400nm Mean, 400nm Kurtosis	0.15
3. Plastic and Metal mines	400nm Max, 900nm Mean, 500nm Max	0.80
4. Metal mines only	400nm Max,500nm Kurtosis, 700nm Energy	0.40
5. Plastic mines only	500nm Mean, 400nm Dev, 500nm Kurtosis	0.03

4.2 Results and discussion

Figs. 4a - d show the resulting labeled and post-processed images of data set t1cfl08. This data set was chosen for an example because it contains a good sampling of metal and plastic mines.

We calculate the Probability of Detection (number of mine regions correctly identified / Total number of mine regions in the scene) and the Probability of False Alarm (number of background pixels identified as targets / Total number of background pixels in the scene) to assess the performance of our analysis. We consider a target correctly identified if the location of the centroid of the region is within a circular neighborhood the size of the radius of the target according to ground truth.

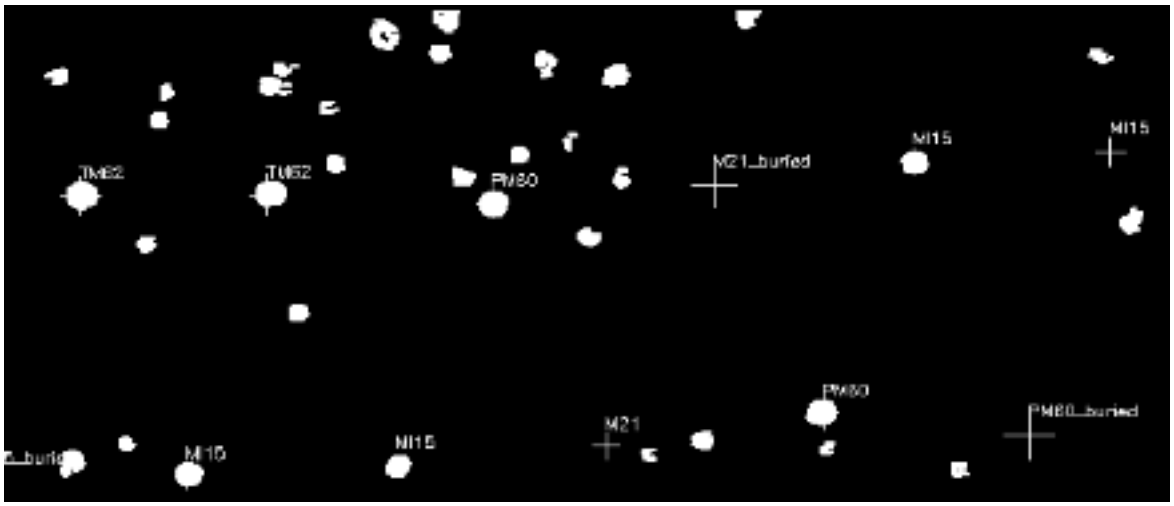

Figure 4a. T1CFL08 post processed: Trained on all targets; PD=1.00, PFA=0.0586

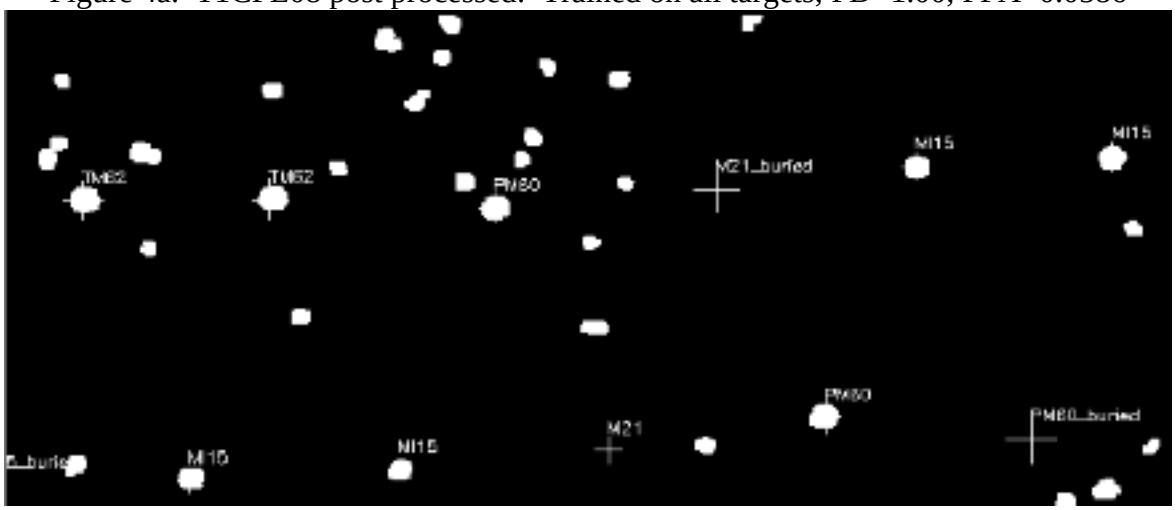

Figure 4b. T1CFL08 post processed: Trained without Resolution Panels; PD=1.00, PFA=0.036

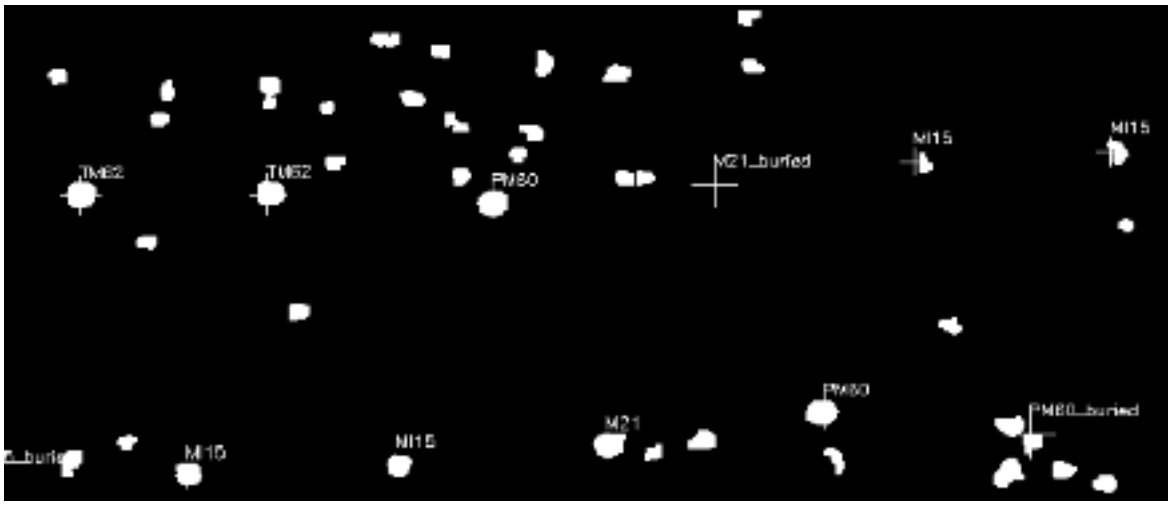

Figure 4c. T1CFL08 post processed: Trained with Metal and Plastic; PD=1.00, PFA=0.017

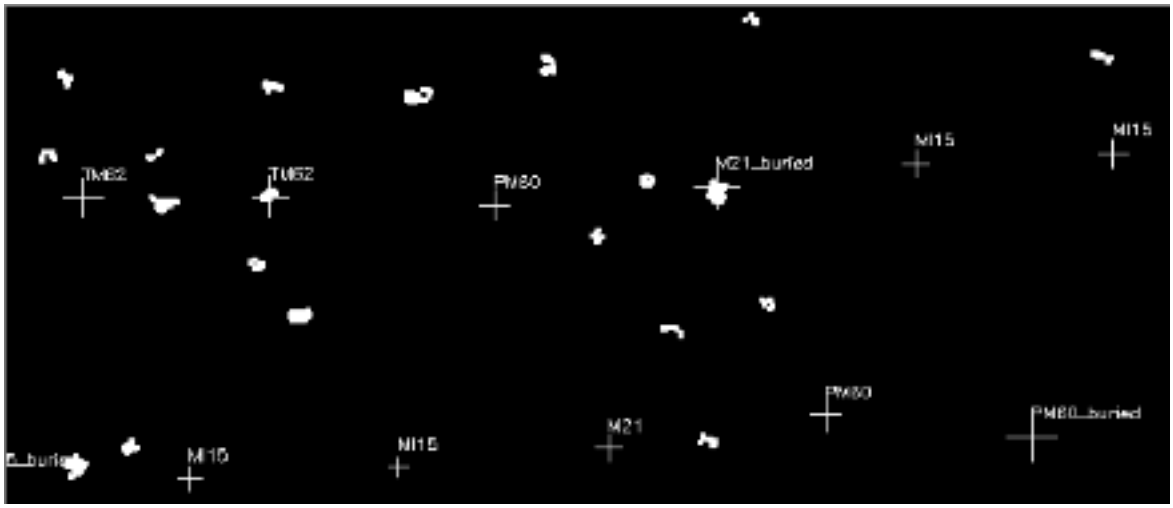

Figure 4d. T1CFL08 post processed: Trained with Plastic; PD =0.11, PFA=0.007

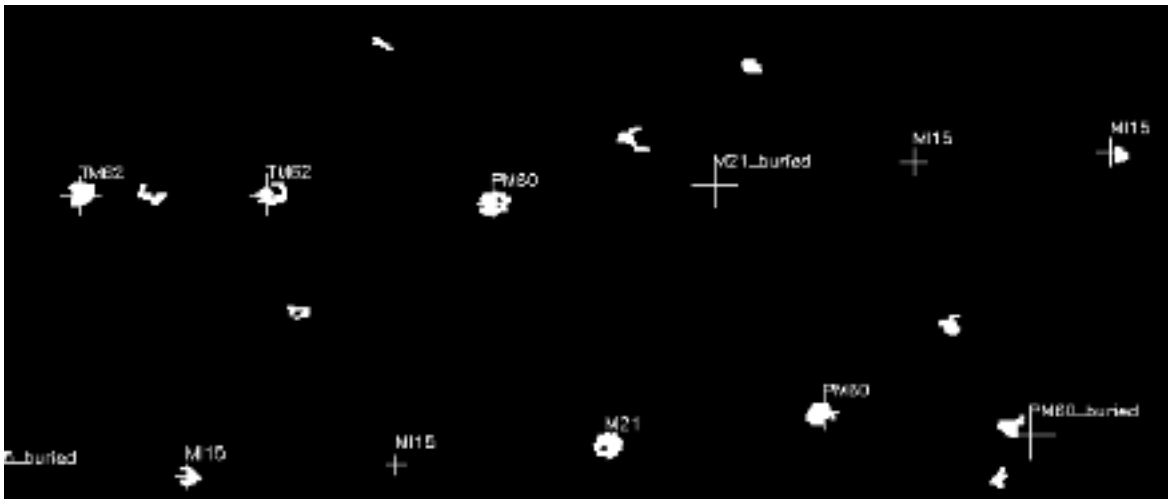

Figure 4e. T1CFL08 post processed: Trained with Metal; PD=0.77, PFA=0.0034

For the case in which all objects in the ground truth were used to train the PNN (fig. 4a) we successfully detected all of the mines so that out of the objects labeled we have a probability of detection, $P_D = 1$. The probability of a false alarm (P_{FA}), however, was quite high at .0586.

Training on all objects except the resolution panels (fig. 2b) gave similar results with $P_D = 1$ and $P_{FA} = .036$. Training on just metal and plastic mines (fig 2c) gave $P_D = 1$ and $P_{FA} = .017$. Finally training on plastic mines only gave poor results (a low P_D), with $P_D = .11$ and $P_{FA} = .0034$ (fig. 4d).

The best results occurred when we used only metal mines in the training tables (fig4e). While P_D in this example dropped to .77, P_{FA} showed a dramatic improvement at .0034. Examples of other data sets (fig 5) show that this trend in performance is consistent throughout the data sets. Data set T1CFL011 has $P_D = 1$ and $P_{FA} = .0013$. T1CFL14 shows excellent results in the labeled image only (fig 5b) with $P_D = 1$ and $P_{FA} = .0028$. However, size constraints applied to this image produced a degradation in the

results (fig 3c), with P_D = .41 and P_{FA} = .00089. Clearly, more investigation needs to be done to find the optimal size constraint that will allow the procedure to be robust with respect to all the data sets.

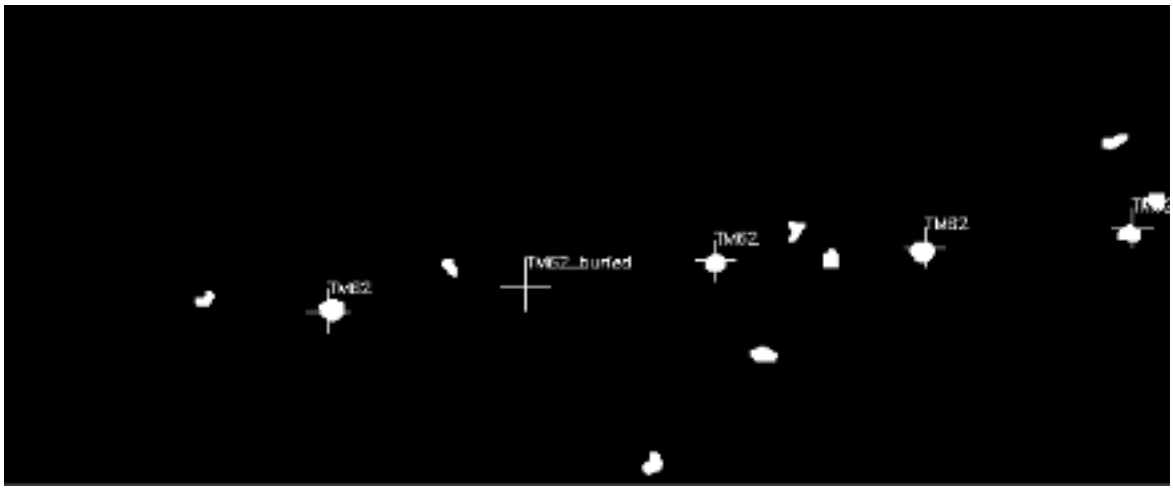

Figure 5a. T1CFL11 post processed . PD=1.0, PFA=0.0030

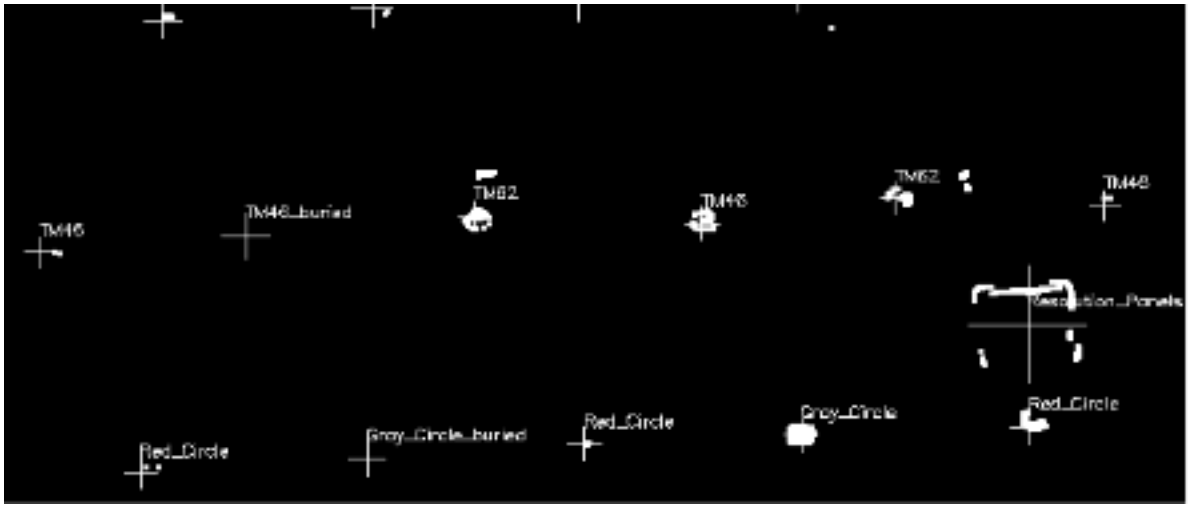

Figure 5b. T1CFL14: Metal only; PD=1.0, PFA=.0028

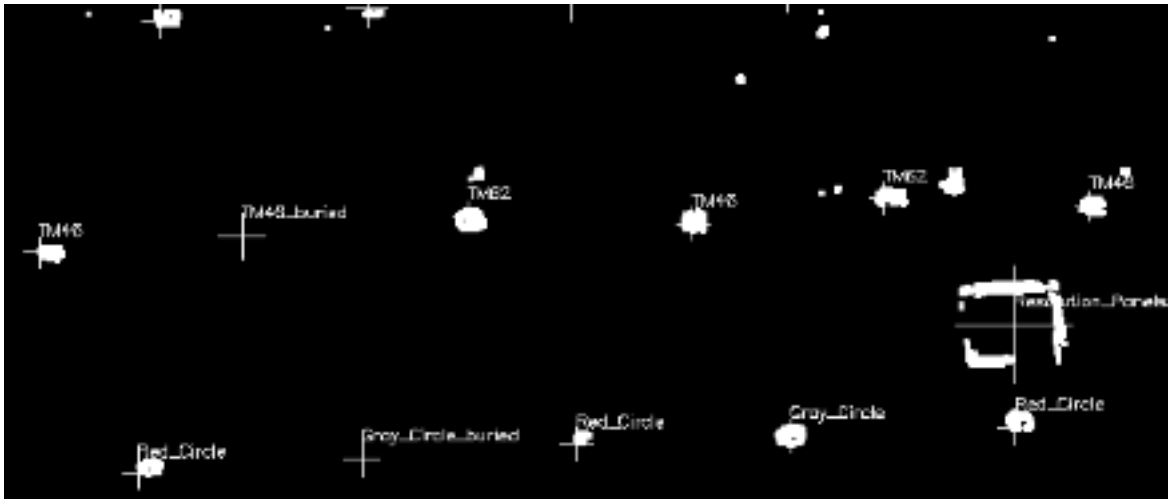

Figure 5c TICFL14: Metal only; Size constraints applied. PD=0.41, PFA=0.00089

The apparent success of the mines-only training set is most likely due to the higher contrast between metal mines and background. When training sets are used containing plastic mines the performance is diminished by the fact that in many of the bands in the data sets the plastic mines have characteristics much like the background. This is evident in fig. 4d.

5. FUTURE WORK

The results of the processing are very encouraging. The probabilities of detection and false alarm achieved during training and testing were sufficient for locating minefields from a qualitative visual inspection of the labeled and postprocessed images. They are very good, given the constraints on the analysis. The current limitations, or barriers to progress, are due primarily to the following factors: (1) The data contain a relatively small number of mine training samples for the individual mine types. Of course, the principles of supervised learning rely on the assumption that the training set is sufficiently large and sufficiently representative of the test set to allow for good detection. This data set is smaller than it should be for good training. (2) The images contain a spatial trend in intensity values that is caused by problems with the camera. We removed the trend with image processing methods, but it should, ideally, be removed by using the hardware. CSS is aware of this problem and corrective action is being taken. (3) One more limitation we encountered with the data was that the resolution of the camera system was often insufficient to distinguish between circles and squares the size of a typical mine. This limits the ability to use features based on shape for object detection. More sophisticated methods that make use of shape information could be used if the camera resolution were greater. If the resolution is fixed for all possible scenarios, this issue may be moot. (4) Only a very simple analysis was attempted, due to funding and time constraints. Further work involving use of more powerful algorithms promises to provide improved results.

Performance improvements are expected if we address these current limitations as follows: (1) Increasing the number of training samples by conducting new experiments, (2) Solving the problems associated with the camera, (spatial trend and resolution limitations), (3) Using shape-based features on higher-resolution images, (4) Conducting a more thorough study of the features and of which ones add the most value to the analysis. To date, we have used only amplitude features and a suboptimal but fast and efficient feature selection algorithm. A better study can be done using our optimal algorithm. We

can also apply more physical knowledge and human judgment to the feature selection process to provide a more thorough analysis, (5)Tuning the tile size for optimum performance.

6. CONCLUSIONS

The algorithms developed at LLNL for buried land mine detection have been applied to the CSS flight test data containing mostly surface land mines, and a few buried mines. The algorithm performance is good and quite encouraging. Probability of detection is very high (often close or equal to one). The probability of false alarm is low enough to allow qualitative detection of a minefield from visual inspection of the labeled and postprocessed images. The main issue of concern is that some labeled images are corrupted by small false alarm regions. Many of these are eliminated by applying size constraints, and we expect that many more could be eliminated by shape constraints. The current barriers to progress are small sample size, relatively low camera resolution, spatial trends in camera intensity, and fiscal constraints. We expect that the performance can be improved significantly by addressing these current limitations in future work.

7. ACKNOWLEDGMENTS

The authors gratefully acknowledge our sponsors; David Vaughn of the Marine Corp., Quantico, VA, Ned Witherspoon and Ab Dubey of CSS, who also provided the image data, and Mark Eckart (LLNL), who served as LLNL co-sponsor. We thank Michael Carter and Charles Anderson of LLNL, who provided valuable input for the analysis.
"Work performed under the auspices of the U.S. Department of Energy by the Lawrence Livermore National Laboratory under contract number W-7405-ENG-48."

8. REFERENCES

1. G. Wolberg, Digital Image Warping, IEEE Computer Society Press, 1990.
2. W.K. Pratt, Digital Image Processing, 2nd Edition, Wiley, pp.559-561.
3. R.M. Welch, K. Kuo, S.K. Sengupta, "Cloud and Surface Textural Features in Polar Regions", IEEE Trans. Geoscience and Remote Sensing, vol. 28, No. 4, pp. 520-528, July 1990.
4. T.Y. Young and K.S. Fu, Handbook of Pattern Recognition and Image Processing, Academic Press, 1986.
5. R.P. Lippmann, "An Introduction to Computing with Neural Nets," IEEE ASSP Magazine, April, 1987, pp. 4-22.
6. D.E. Specht, "Probabilistic Neural Networks," Neural Networks, Vol. 3, pp.109-118, 1990.
7. F. Roeske, W. Aimonetti, G. Clark, J. Donetti, D. Fields, P. Schaich, S. K. Sengupta, R. J. Sherwood, "Preliminary Results of LLNL Analysis of NCSS Cobra Flight Test Data for Mine Detection," Lawrence Livermore National Laboratory report UCRL- , September, 1994 (In Preparation).
8. G. A. Clark, S. K. Sengupta, P. C. Schaich, R. J. Sherwood, M. R. Buhl, J. E. Hernandez, D. J. Fields, and M. R. Carter, "Data Fusion for the Detection of Buried Land Mines," International Symposium on Substance Identification Technologies, European Optical Society (EOS)/The International Society for Optical Engineering (SPIE), Congress Centre Innsbruck, Innsbruck, Austria, October 4-8, 1993.

9. G. A. Clark, Sailes K. Sengupta, Robert J. Sherwood, Jose E. Hernandez, Michael R. Buhl, Paul C. Schaich, Ronald J. Kane, Marvin J. Barth and Nancy K. DelGrande, "Sensor Feature Fusion for Detecting Buried Objects," SPIE's 1993 International Symposium and Exhibition on Optical Engineering and Photonics in Aerospace and Remote Sensing, Conference on Underground and Obscured Imaging and Detection, Orlando, FLA, April 12-16, 1993.
10. Michael R. Buhl, Jose E. Hernandez, Gregory A. Clark, and Sailes K. Sengupta, "Dual-Band Infrared Buried Mine Detection Using a Statistical Pattern Recognition Approach," Lawrence Livermore National Laboratory report UCRL-ID-114838, August, 1993.
11. G. A. Clark, S. K. Sengupta, M. R. Buhl, R. J. Sherwood, P. C. Schaich, N. Bull, R. J. Kane, and M. J. Barth, "Detecting Buried Objects by Fusing Dual-Band Infrared Images," *Invited paper*, Asilomar Conference on Signals, Systems and Computers, Pacific Grove, CA, Nov. 1-3, 1993.
12. N. K. DelGrande, P. F. Durbin, D. E. Perkins, G. A. Clark, R. J. Sherwood, J. E. Hernandez, and A. B. Shapiro, "Dual-Band Infrared Capabilities for Imaging Buried Object Sites, " SPIE's 1993 International Symposium and Exhibition on Optical Engineering and Photonics in Aerospace and Remote Sensing, Conference on Underground and Obscured Imaging and Detection, Orlando, Florida, April 12-16, 1993.
13. M. R. Buhl, J. E. Hernandez, S. K. Sengupta, G. A. Clark, "Dual-Band Infrared Buried Mine Detection Using a Statistical Pattern Recognition Approach," Lawrence Livermore National Laboratory report UCRL-114838, August, 1993.
14. G. A. Clark, "Report on the Workshop on Automatic Target Recognition (ATR) for Airborne Standoff Mine Detection (ASMD)," Sponsored by DARPA and the U. S. Marines at the System Planning Corporation, Arlington, VA (Workshop: December 9-10, 1992), March 5, 1993.
15. G. A. Clark (Workshop Chairman), J. E. Hernandez, S. K. Sengupta, R. J. Sherwood, M. R. Buhl, P. C. Schaich, R. J. Kane, M. J. Barth, and N. K. DelGrande, "Computer Vision and Sensor Fusion for Detecting Buried Mines," Workshop on Automatic Target Recognition (ATR) for Airborne Standoff Mine Detection (ASMD), Sponsored by DARPA and the U. S. Marines at the System Planning Corporation, Arlington, VA, December 9-10, 1992.
16. G. A. Clark, J. E. Hernandez, S. K. Sengupta, R. J. Sherwood, P. C. Schaich, M.R. Buhl, R. J. Kane, M.J. Barth, and N. K. DelGrande, "Computer Vision and Sensor Fusion for Detecting Buried Objects," Asilomar Conference on Signals, Systems and Computers, Pacific Grove, CA, October 26-28, 1992.
17. N.K. Del Grande, G.A. Clark, P.F. Durbin, D.J. Fields, J.E. Hernandez, and R.J. Sherwood, "Buried Object Remote Detection Technology for Law Enforcement," SPIE Orlando '91 Symposium, Orlando, Florida, April 1-5, 1991.
18. G.A. Clark, J.E. Hernandez, N.K. Del Grande, R.J. Sherwood, S-Y Lu, and P.F. Durbin, "Computer Vision for Locating Buried Objects," Twenty-Fifth Annual Asilomar Conference on Signals, Systems, and Computers, Pacific Grove, California, November 4-6, 1991.
19. Airborne Detection of Buried Minefields, Energy and Technology Review, Lawrence Livermore National Laboratory, December, 1991.
20. Hu, M.K., 1962. "Visual Pattern Recognition by Moment Invariants", IRE Trans. Info. Theory, vol. IT-8, pp. 179-187.
21. Devijver, P. A. and J. Kittler, 1982. Pattern Recognition: A Statistical Approach. Prentice Hall, Englewood Cliffs, 448 pp.
22. Narendra, P. and K. Fukunuga, 1977. "A Branch and Bound Algorithm for Feature Subset Selection", IEEE Trans. Comp., vol. C-26, no. 9, Sept. 1977, pp 917-922.
23. Specht, D., 1990b. "PNN and Polynomial Adaline as Complementary Techniques for Classification", IEEE Trans. Neural Networks, vol. 1, pp 111-121.
24. Cacoullos, T., 1966. "Estimation of a Multivariate Density", Annals of the Institute of Statistical Mathematics (Tokyo), vol. 18(2), pp. 179-198.

25. Murthy, V.K. 1965. “Estimation of Probability Density”, Annals of Mathematical Statistics, vol. 36, pp. 1027-1031.
26. Murthy, V. K., 1966. Nonparametric Estimation of Multivariate Densities with Applications, in P.R. Krishnaiah (Ed.), Multivariate Analysis (pp. 43-58), New York: Academic Press.
27. Parzen, E., 1962. “On Estimation of a Probability Density Function and Mode”, Annals of Mathematical Statistics, vol. 33, pp. 1065-1076.
28. Rumelhart, D.E., J.L. McClelland and the PDP Research Group, 1986. Parallel Distributed Processing: Explorations in the Microstructure of Cognition, I & II, Cambridge, MA: MIT Pr. 611 pp.
29. Hand, D. J., 1981. Discrimination and Classification, New York, L. Wiley and Sons, pp. 218.
30. A. K. Jain, Fundamentals of Digital Image Processing, Prentice Hall, 1989.
31. Gonzales, Rafael C. and Wintz, Paul., Digital Image Processing, second edition. Addison-Wesley Publishing Company, 1987.
32. Hall, Ernest L., Computer Image Processing and Recognition. Academic Press, 1979.
33. Rosenfeld, Azriel and Kak, Avinash C, Digital Picture Processing., Academic Press, 1976.
34. R. M. Haralick and L. G. Shapiro, *Computer and Robot Vision*, Vol. 1, Addison Wesley, 1992.

Technical Information Department • **Lawrence Livermore National Laboratory**
University of California • Livermore, California 94551